

AMENDMENTS TO THE CLAIMS

1. (Currently amended) A caching device to operate as an intermediary node on a network, the device comprising:

a cache to store content requestable by a client on the network;

a user interface to enable a user to specify a set of forwarding rules for forwarding requests on the network;

a database to store the set of forwarding rules;

a request processing unit to receive a request from the client;

a rule evaluator to evaluate the set of forwarding rules to identify within the set of forwarding rules a rule which applies to the request, such that the request processing unit attempts to forward the request to a destination selected according to said rule, wherein the rule indicates a host in a defined forwarding hierarchy; and

a rule engine to determine an availability of the host indicated in said rule and, if the host is available, to select the host as a forwarding destination and to cause the request processing unit to forward the request to the host according to said rule, wherein the host is determined to be available if the host responds to a connection attempt within a timeout period and, if more than one host is available as indicated by said rule, to select the host based on a delivery factor included in said rule, and

wherein, if the request processor is unable to forward the request according to said rule, the rule evaluator resumes evaluating the set of forwarding rules to identify another rule corresponding to the request.

2. (Previously amended) A device as recited in claim 1, wherein the user may specify one or more of the forwarding rules to indicate a manner of forwarding the request within the forwarding hierarchy.
3. (Previously amended) A device as recited in claim 1, wherein the forwarding hierarchy is a cache hierarchy.
4. (Original) A device as recited in claim 3, further comprising a cache to store content requestable by a client on the network, wherein the request processing unit forwards the request only in the event of a cache miss, wherein the request is a request for content on the network.
5. (Canceled)
6. (Original) A device as recited in claim 1, wherein the rule evaluator identifies the rule which applies to the request by determining that a condition in the rule involving a variable in the request is satisfied.
7. (Canceled)
8. (Currently amended) A device as recited in claim 1, wherein said rule comprises a plurality of destinations, and wherein the rule engine selects a destination from among the plurality of destinations as a forwarding destination for the request, based on at least one ~~the~~ delivery factor included in the rule.
9. (Currently amended) A device as recited in claim 8~~1~~, wherein the ~~at least one~~ delivery factor comprises a specified distribution method for the request.
10. (Currently amended) A device as recited in claim 8~~1~~, wherein the ~~at least one~~ delivery factor comprises an indication of a current load ~~on at least one of the~~ destinations.

11. (Currently amended) A device as recited in claim 8~~1~~, wherein the ~~at least one~~ delivery factor comprises a weighting of the ~~plurality of destinations~~ indicating a preferred distribution of forwarding requests ~~between the plurality of destinations~~.

12. (Original) A device as recited in claim 1, wherein the user interface is further to enable the user to modify the set of forwarding rules.

13. (Original) A device as recited in claim 1, wherein the user interface is further to enable the user to specify a sequence in which the rules of the set of forwarding rules are evaluated in response to the request.

14. (Original) A device as recited in claim 1, wherein the user interface is further to enable the user to selectively enable or disable a default forwarding rule.

15. (Currently amended) A device to operate as an intermediary node on a network, the device comprising:

a processor;

a network interface to allow the device to communicate on the network; and

a storage facility to store program code for execution by the processor to cause the device to provide a user interface to enable a user to specify a set of forwarding rules, receive a request for content from a client, determine whether the content is cached locally in said device, and in response to a determination that the content is not cached locally in the device, evaluate the set of forwarding rules to identify a rule in the set of forwarding rules which should be applied to the request, determine an availability of a host indicated in said rule and select the host as a forwarding destination for the request if the host is available, wherein the host is determined to be available if the host responds to a connection attempt within a timeout period, select the host based on a delivery factor included in said rule if more than one host is available as indicated by said rule, forward the request on the network according to said rule, and resume

evaluating the set of forwarding rules to identify another rule corresponding to the request if the device is unable to forward the request according to said rule.

16. (Original) A device as recited in claim 15, wherein the device operates within a defined forwarding hierarchy, and the user may specify one or more of the forwarding rules to indicate a manner of forwarding the request within the forwarding hierarchy.

17. (Original) A device as recited in claim 16, wherein the forwarding hierarchy is a cache hierarchy.

18. (Original) A device as recited in claim 17, further comprising a cache to store content requestable by a client on the network, wherein the device forwards the request only in the event of a cache miss at the device.

19. (Original) A device as recited in claim 15, wherein the request is a request for content on the network.

20. (Canceled)

21. (Previously amended) A device as recited in claim 15, wherein the program code is further to cause the device, when executed by the processor, to identify the rule which should be applied in response to the request by identifying a correspondence between a variable in the request and a variable in the rule.

22. (Original) A device as recited in claim 15, wherein the user interface is further to enable the user to modify the set of forwarding rules.

23. (Currently amended) An intermediary network node comprising:

means for receiving a request for content on a network;

means for determining a forwarding destination for the request in a defined forwarding hierarchy, by applying a set of user-specified forwarding rules to the request;

means for determining an availability of a host indicated in said rule and selecting the host as a forwarding destination for the request if the host is available, wherein the host is determined to be available if the host responds to a connection attempt within a timeout period, and when more than one host is available as indicated by said rule, the host is selected based on a delivery factor included in said rule;

means for forwarding the request according to the determined forwarding destination; and

means for resuming evaluation of the set of forwarding rules to identify another rule corresponding to the request if said forwarding the request is unsuccessful.

24. (Original) An intermediary network node as recited in claim 23, wherein the forwarding hierarchy is a cache hierarchy.

25. (Original) An intermediary network node as recited in claim 24, wherein the intermediary node is a network caching device, and wherein said determining and said forwarding are performed only if the request results in a cache miss at the intermediary network node.

26. (Currently amended) A caching device to operate within a cache hierarchy on a network, the caching device comprising:

a cache to store content requestable by a client on the network;

a request processing unit to receive a request for content from the client, and to forward the request on the network based on a set of forwarding rules in the event of a cache miss;

a user interface to enable a user to specify the set of forwarding rules, such that the user may specify one or more forwarding rules to indicate a host in the cache hierarchy as a destination for a corresponding request;

a database to store the set of forwarding rules;

a rule evaluator to evaluate the set of forwarding rules in response to the cache miss, to identify a rule in the set of forwarding rules which applies to the request; and

a rule engine to determine an availability of a host indicated in the rule, and to select the host as a forwarding destination for the request if the host is available, wherein the host is determined to be available if the host responds to a connection attempt within a timeout period, and if more than one host is available as indicated by the rule, to select the host based on a delivery factor included in the rule, the rule engine further to indicate to the request processing unit if the host is available to cause the request processing unit to forward the request to the host, wherein, if the request processing unit is unable to forward the request according to said rule, the rule evaluator resumes evaluating the set of forwarding rules to identify another rule corresponding to the request.

27. (Original) An intermediary network node as recited in claim 26, wherein the rule evaluator identifies the rule which applies to the request by determining that a condition in the rule is satisfied.

28. (Original) An intermediary network node as recited in claim 27, wherein the condition involves a variable in the request.

29. (Original) An intermediary network node as recited in claim 26, wherein said rule comprises a plurality of destinations, and wherein the rule engine selects said destination from among the plurality of destinations, based on the availability of each of the plurality of destinations.

30. (Original) An intermediary network node as recited in claim 26, wherein the user interface is further to enable the user to modify the set of forwarding rules.

31. (Original) An intermediary network node as recited in claim 26, further comprising a rule encoder to encode into a uniform syntax forwarding rules specified by the user.

32. (Currently amended) A network caching device to operate within a defined cache hierarchy on a network, the caching device comprising:

a cache to store content from an origin server on the network;

an application to receive a request for content from a client via the network, and to forward the request on the network based on a set of forwarding rules in the event of a cache miss;

a user interface to enable a user to specify and modify the set of forwarding rules;

a rule encoder to encode into a uniform syntax forwarding rules specified by the user;

a rules database to store the encoded forwarding rules;

a rule evaluator to evaluate the set of forwarding rules sequentially in response to the cache miss, to identify a rule in the set of forwarding rules which applies to the request, by identifying a correspondence between a variable in the request and a variable in the rule, the rule specifying a host within the cache hierarchy as a forwarding destination for the request; and

a rule engine to determine an availability of the host and to select the host as said forwarding destination for the request if the host is available, the rule engine further to indicate the host to the application layer if the host is available, wherein the host is determined to be available if the host responds to a connection attempt within a timeout period, to cause the application layer to attempt to establish a connection with the host, and if more than one host is available as indicated by said rule, to select the host based

on a delivery factor included in said rule, such that the application layer forwards the request to the host upon successfully establishing the connection,

wherein if the application layer is unable to forward the request according to the rule, the rule evaluator resumes evaluating the set of forwarding rules to identify another rule corresponding to the request.

33. (Original) A network caching device as recited in claim 32, wherein said rule comprises a plurality of destinations, and wherein the rule engine selects said destination from among the plurality of destinations, based on the availability of each of the plurality of destinations.

34. (Currently amended) A method comprising:

receiving, at an intermediary network node, a request for content on a network;

determining, in the intermediary network node, a forwarding destination in a defined forwarding hierarchy, by applying a set of user-specified forwarding rules to the request;

determining an availability of the forwarding destination wherein the forwarding destination is determined to be available if the forwarding destination responds to a connection attempt within a timeout period, and if more than one forwarding destination is determined to be available as indicated by the forwarding rule, selecting the forwarding destination based on a delivery factor included in the forwarding rule;

forwarding the request according to the determined forwarding destination if the forwarding destination is available; and

resuming evaluation of the set of forwarding rules if said forwarding the request is unsuccessful.

35. (Original) A method as recited in claim 34, wherein the forwarding hierarchy is a cache hierarchy.

36. (Original) A method as recited in claim 35, wherein the intermediary node is a network caching device, and wherein said determining and said forwarding are performed only if the request results in a cache miss at the intermediary network node.

37. (Currently amended) A method of operating a caching device in a cache hierarchy on a network, the method comprising:

 caching content on the network;

 providing a user interface to enable a user to specify a set of forwarding rules;

 storing the set of forwarding rules;

 receiving a request from the client;

 evaluating the set of forwarding rules if the request produces a cache miss, to identify a rule in the set of forwarding rules that applies to the request;

 determining an availability of a host indicated in the rule, wherein the host is determined to be available if the host responds to a connection attempt within a timeout period, and if more than one host is available as indicated by the rule, selecting the host based on a delivery factor included in the rule;

 attempting to establish a connection to the host if the host is available;

 forwarding the request to the host; and

 resuming said evaluating to identify another rule having a correspondence to the request if said attempting to establish the connection is unsuccessful.

38. (Canceled)

39. (Original) A method as recited in claim 37, further comprising, if said attempting to establish the connection is unsuccessful:

determining whether a second available host is indicated in the rule, and if so, attempting to establish a connection to the second available host; and

forwarding the request to the second available host.

40. (Original) A method as recited in claim 37, wherein said providing the user interface comprises providing the user interface such that the user may specify one or more forwarding rules to indicate a host in the cache hierarchy as a destination for a corresponding request.

41. (Original) A method as recited in claim 37, wherein the rule is determined to apply to the request if a condition in the rule is satisfied by the request.

42. (Original) A method as recited in claim 41, wherein the condition is a function of a variable in the request.

43. (Original) A method as recited in claim 37, wherein said providing the user interface comprises providing the user interface to enable the user to modify the set of forwarding rules.

44. (Previously presented) A device as recited in claim 1, wherein the request processing unit is further to use a timeout period in attempting to establish a connection with the forwarding destination, the timeout period having been computed based on information indicative of a responsiveness of the forwarding destination.

45. (Previously presented) A device as recited in claim 44, wherein the information indicative of the responsiveness of the forwarding destination comprises information indicative of a loading on, or a response time of, the forwarding destination.

46. (Previously presented) A device as recited in claim 44, wherein the timeout period is computed in response to the request being received.

47. (Previously presented) A method as recited in claim 34, further comprising, prior to said forwarding the request:

acquiring information about the forwarding destination, the information being indicative of a responsiveness of the forwarding destination;

computing a timeout period based on the information indicative of the responsiveness of the forwarding destination; and

using the computed timeout period, in response to the request, in attempting to establish a connection with the forwarding destination.

48. (Previously presented) A method as recited in claim 47, wherein said computing a timeout period is done in response to receiving the request.

49. (Currently amended) A proxy cache to operate as an intermediary node on a network, the proxy cache comprising:

a user interface to enable a user to specify a set of forwarding rules for forwarding requests on the network, wherein ~~at least one~~ a rule in said set of forwarding rules can specify a plurality of destinations;

a database to store the set of forwarding rules;

a request processing unit to receive a request at the proxy cache from a client;

a rule evaluator to evaluate the set of forwarding rules to identify a rule in the set of forwarding rules which applies to the request; and

a rule engine to select a destination from among the plurality of destinations based on ~~at least one~~ a delivery factor included in the rule, the ~~at least one~~ delivery factor comprising ~~at least one~~ of: a specified distribution method for the request, an indication of a current load on ~~at least one~~ of the plurality of destinations, or a weighting of the plurality of destinations indicating a preferred distribution of forwarding requests between the plurality of destinations, the rule engine further to cause the request processing unit to forward the request to the destination if the destination is available, wherein the destination is determined to be available if the destination responds to a connection attempt within a timeout period, and if more than one destination is available, the rule engine to further select one of the available destinations based on one of the delivery factors.

50. (Canceled)

51. (Previously amended) A device as recited in claim 49, wherein the proxy cache is configured to operate within a defined cache hierarchy, and the user may specify one or more of the forwarding rules to indicate a manner of forwarding the request within the cache hierarchy.

52. (Previously presented) A device as recited in claim 49, further comprising a cache to store content requestable by a client on the network, wherein the request processing unit forwards the request only in the event of a cache miss, wherein the request is a request for content on the network.

53-56. (Canceled)

57. (Previously presented) A device as recited in claim 49, wherein the user interface is further to enable the user to modify the set of forwarding rules, to specify a sequence in which the forwarding rules are to be evaluated, and to selectively enable or disable a default forwarding rule.

58. (Currently Amended) A method of operating a network cache, comprising:

receiving from a client a request for content;

in response to a cache miss caused by the request, identifying a rule corresponding to the request from a set of rules;

determining availability of a host indicated in the rule, wherein the host is determined to be available if the host responds to a connection attempt within a timeout period, and if more than one host is available as indicated by the rule, selecting the host based on a delivery factor included in the rule;

selecting the host as a forwarding destination in response to determining that the host is available;

determining whether the request is successfully forwarded to the host; and

identifying another rule corresponding to the request if the request is not successfully forwarded to the host.

59. (New) A device as recited in claim 1, wherein the delivery factor comprises an indication of a link bandwidth between the request processing unit and the destinations.

60. (New) A device as recited in claim 15, wherein the delivery factor comprises at least one of:

a specified distribution method for the request;

an indication of a current load;

a weighting indicating a preferred distribution of forwarding requests; or

an indication of a link bandwidth.

61. (New) An intermediary node as recited in claim 23, wherein the delivery factor comprises at least one of:

a specified distribution method for the request;

an indication of a current load;

a weighting indicating a preferred distribution of forwarding requests; or

an indication of a link bandwidth.

62. (New) A caching device as recited in claim 26, wherein the delivery factor comprises at least one of:

a specified distribution method for the request;

an indication of a current load;

a weighting indicating a preferred distribution of forwarding requests; or

an indication of a link bandwidth.

63. (New) A network caching device as recited in claim 32, wherein the delivery factor comprises at least one of:

a specified distribution method for the request;

an indication of a current load;

a weighting indicating a preferred distribution of forwarding requests; or

an indication of a link bandwidth.

64. (New) A method as recited in claim 34, wherein the delivery factor comprises at least one of:

a specified distribution method for the request;

an indication of a current load;

a weighting indicating a preferred distribution of forwarding requests; or

an indication of a link bandwidth.

65. (New) A method as recited in claim 37, wherein the delivery factor comprises at least one of:

a specified distribution method for the request;

an indication of a current load;

a weighting indicating a preferred distribution of forwarding requests; or

an indication of a link bandwidth.

66. (New) A method as recited in claim 58, wherein the delivery factor comprises at least one of:

a specified distribution method for the request;

an indication of a current load;

a weighting indicating a preferred distribution of forwarding requests; or

an indication of a link bandwidth.